

Light and Shading

EECS 442 – Prof. David Fouhey
Winter 2019, University of Michigan

http://web.eecs.umich.edu/~fouhey/teaching/EECS442_W19/



Administrivia

- I sent out requests for waitlist additions. I will continue to add as spots free up but chances are diminishing.
- **PLEASE SIGN UP ON PIAZZA.** There are no secrets on canvas.
- HW1 out. **Any general questions (not about content)?**
- Discussion on Wednesday: image processing / numpy. Materials out on piazza.

Recap: Projection


$$\text{Image} \rightarrow \mathbf{P} = \mathbf{K}[\mathbf{R}, \mathbf{t}] \mathbf{X} \leftarrow \text{World}$$

Intrinsic *Extrinsic*



Recap: Lenses

Pinhole Model



Mathematically correct
Not quite correct in practice
Reasonable approximation

Reality: Lenses

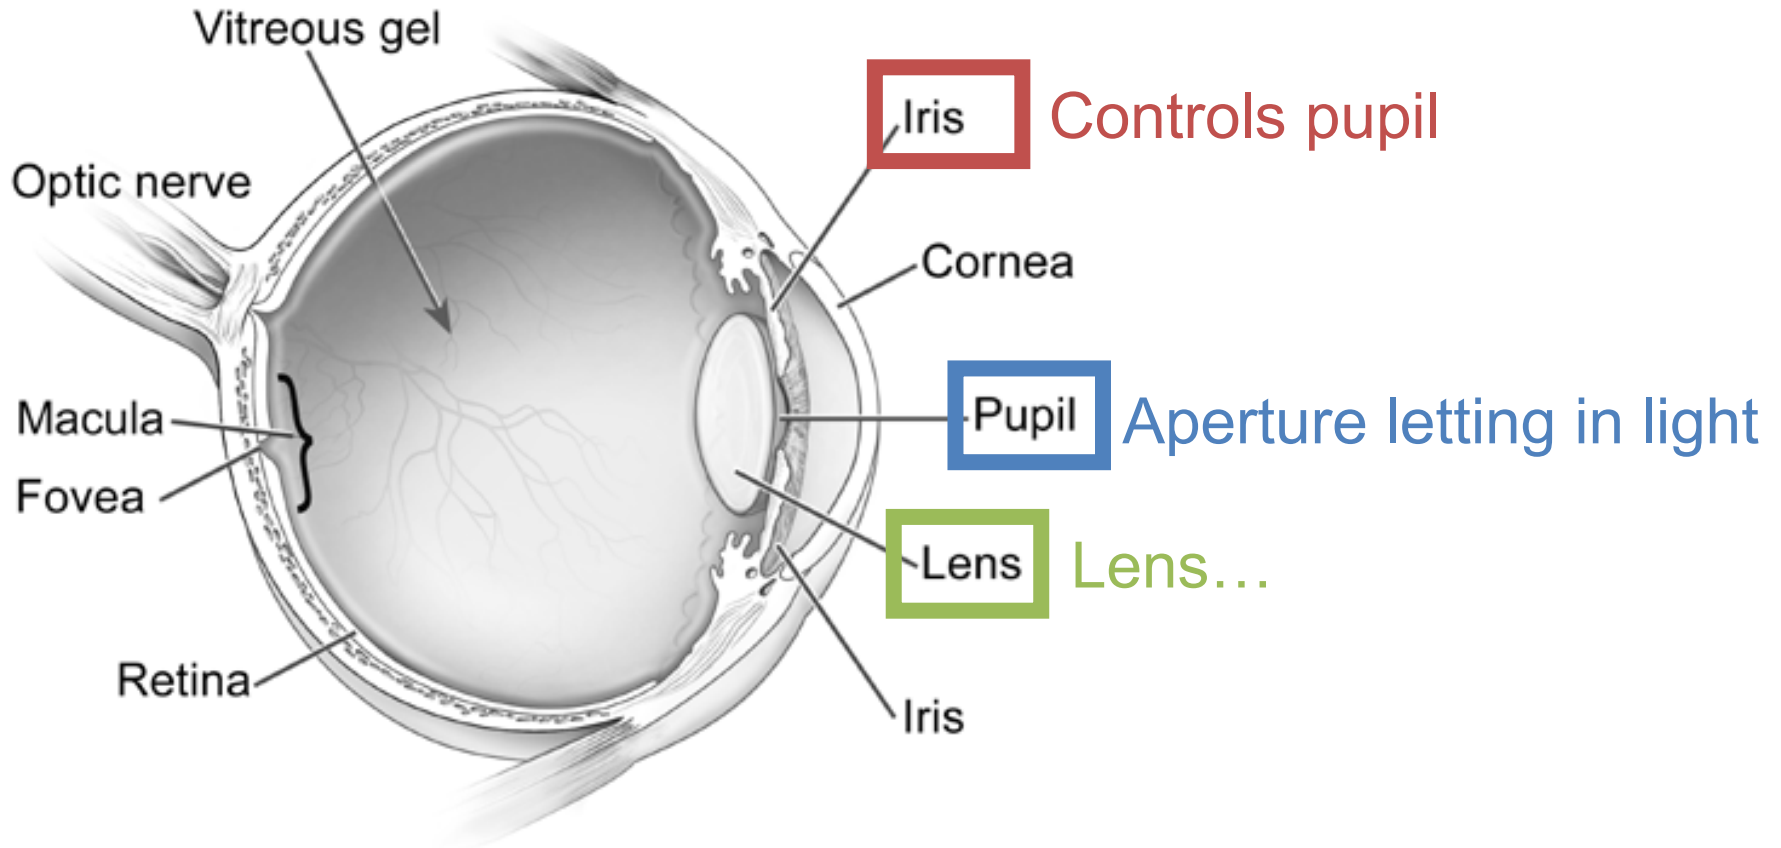


Necessary in practice
Introduce complications
Complications fixable

Today


- A little bit about light and how you represent it
- A little bit about lighting and how it works

Your Very Own Camera



Where's the film/CCD?


Your Very Own Camera



Where's the **film/CCD**?

Demo Time

What is Retina/Film Made Of?




Two Type of Photo Receptors

Cones


cone-shaped
less sensitive
operate in high light
color vision

Rods

rod-shaped
highly sensitive
operate at night
gray-scale vision




Rod / Cone Sensitivity




Rod/Cone Distribution


(a)



(b)

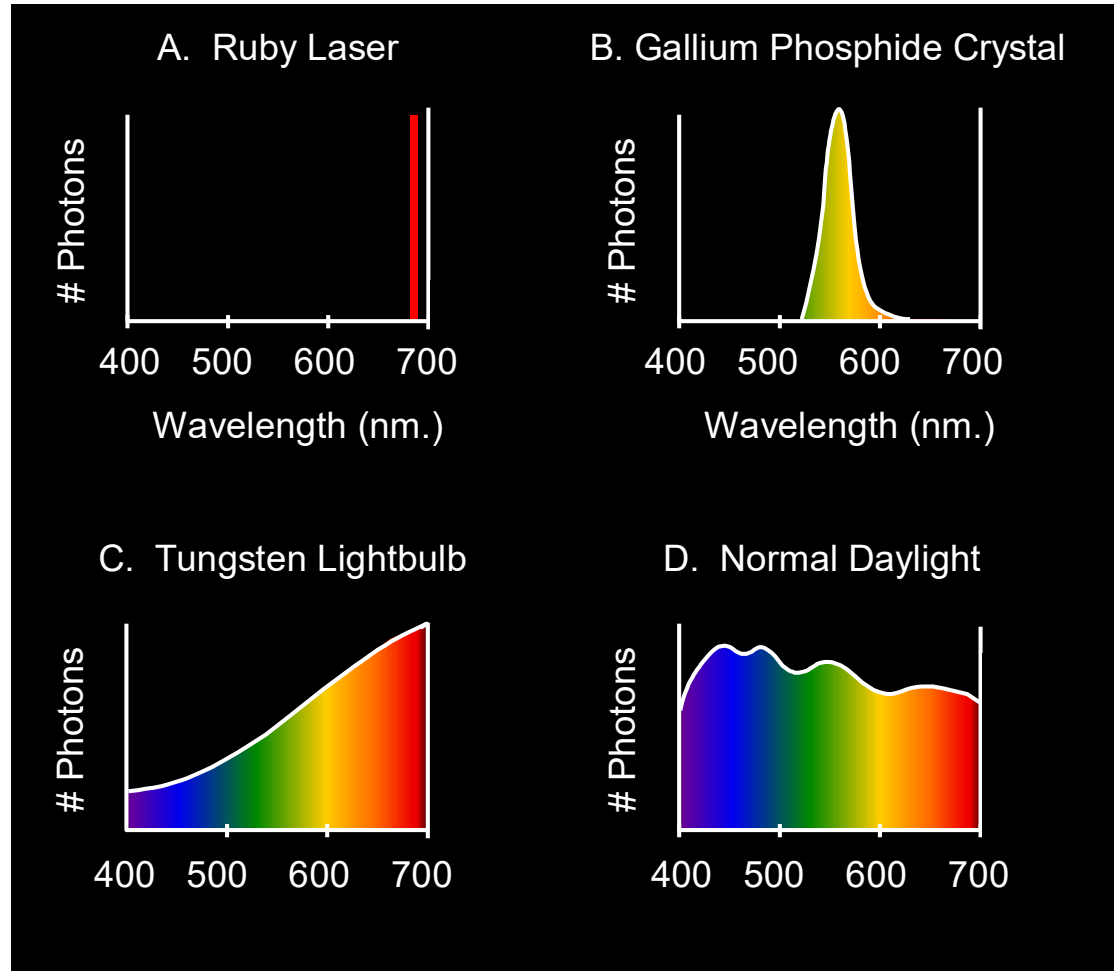


Electromagnetic Spectrum







Why do we see light in these wavelengths?



The Physics of Light



The Physics of Light





The Physics of Light




How Do We Get Light?

Artificial Cones



Estimate RGB
at 'G' cells from
neighboring
values


Color Image



copyright 2000 philg@mit.edu

Color Image

Combined



Red




Green



Blue




Images in Python



5 Things To Always Remember

1. Origin is top left
2. Rows are first index (**what's the fastest direction for accessing?**)
3. Usually referred to as Height x Width
4. Typically stored as uint8 [0,255]
5. for y in range(H): for x in range(W): will run 1 million times for a 1000x1000 image. *A 4GHz processor can do only 4K clock cycles per pixel per second.*

Representing Colored Light



Discussion time: how many numbers do you actually need for colored light? Assume all tuples (R,G,B) are legitimate colors (they are).


One Option: RGB

Pros

1. Simple
2. Common

Cons

1. Distances don't make sense
2. Correlated



R



G



B

RGB





Another Option: HSV

Pros

1. Intuitive for picking colors
2. Sort of common
3. Fast to convert

Cons

1. Not as good as other better spaces



H
(S=1,V=1)



S
(H=1,V=1)



V
(H=1,S=0)

HSV




Photo credit: J. Hays

copyright 2000 phi1ga@vt.edu

Another Option: YCbCr/YUV


Pros

1. Great for transmission / compression


Cons

1. Not as good as other better smart color spaces

Y = 0



Y = 0.5



Y
(Cb=0.5,
Cr=0.5)



Cb
(Y=0.5,
Cr=0.5)



Cr
(Y=0.5,
Cb=0.5)

YCbCr



Photo credit: J. Hays

copyright 2009 philip@mit.edu


Another Option: Lab

Pros

1. Distances correspond with human judgment
2. Safe

Cons

1. Complex to calculate (don't write it yourself, lots of fp calculations)



L
(a=0,b=0)



a
(L=65,b=0)



b
(L=65,a=0)

Lab



Photo credit: J. Hays


Why Are There So Many?

- Each serves different functions
 - RGB: sort of intuitive, standard, everywhere
 - HSV: good for picking, fast to compute
 - YCbCr/YUV: fast to compute
 - Lab: the right(?) thing to do, but “slow” to compute
- Pick based on what you need and don't sweat it: color really isn't crucial


So Far

How do we represent **light**
and its storage on **film**?

Photo. Material





Now




How does the scene
cause that **light**?

Photo. Material




Light and Surfaces



What happens when
light hits a surface?

Light and Surfaces




What happens when light hits a surface?

1. Absorbed

It's absorbed and converted into some other form of energy (e.g., a black shirt getting hot in the sun)

Light and Surfaces




What happens when light hits a surface?

2. Transmitted

Possibly bouncing around before going through or out (e.g. lenses bend and go through, milk bounces around)

Light and Surfaces




What happens when light hits a surface?

3. Reflected

It's reflected back, in one or more directions with varying amounts (e.g., mirror, or a white surface)

Light and Surfaces




What happens when light hits a surface?

4. Everything

All of the above! Real surfaces often have combinations of all of these options.

Modeling Light and Surfaces



Opaque Reflections

Bi-directional reflectance function: % reflected given *incident angle* to light *reflected angle* to the viewer.

Note: have not specified form of function.

Specular and Diffuse Reflection


Same lighting, as close as possible camera settings, but different **location**



Specular and Diffuse Reflection

Diffuse

Specular




Basically same

Totally different



Diffuse Reflection




Lambertian Surface

Light depends **only** on orientation of surface

$$\phi_i, \theta_i$$

to light. Result of random small facets. Looks identical at all views.

Diffuse Reflection



Lambert's Law

N: surface normal


S: source direction **and**
strength

ρ : how much is reflected

$$B = \rho N \cdot S$$

$$B = \rho \|S\| \cos(\theta)$$


Specular Reflection



Specular Surface

Light reflected like a mirror, but spreads out in a “lobe” around the reflection ray

Specular Reflection



Phong Model


V: angle to viewer

R: reflection ray

α : shininess constant

$$B = (V^T R)^\alpha$$

BRDFs can be incredibly complicated...



What Can This Be Used For

Shape from Shading

Lambert's Law: for every pixel i

$$B_i = \rho N_i \cdot S$$

Reflected
Light
(1 dim)

Surface
Orientation
(3? dim)

Illumination
Global,
(3 dim)

Given: illumination and light, recover normals

Potential problems?

Shape From Shading

$$B_i = \rho N_i \cdot S$$

1D, **fixed** actually 2D 3D, **fixed**
unknown

- System of equations that's underdetermined (N equations, 2N unknowns, N+3 known)
- **Solution:** Add more equations that enforce smoothness or finding a single surface.


Realistic Shape From Shading

$$B_i = \rho N_i \cdot S$$

1D, fixed 2D unknown 3D, unknown

- System of equations that's underdetermined (N equations, 2N+3 unknowns)
- **Solution:** need prior beliefs to disambiguate.

Ambiguity



Ambiguity

Humans assume light from above (and the blueness also tells you distance)



Shape from Shading in Practive



<https://www.youtube.com/watch?v=4GiLAOtjHNo>